Is Hamming distance the only way for matching binary image feature descriptors?

E. Bostanci

Brute force matching of binary image feature descriptors is conventionally performed using the Hamming distance. This paper assesses the use of alternative metrics in order to see whether they can produce feature correspondences that yield more accurate homography matrices. Two statistical tests, namely ANOVA (Analysis of Variance) and McNemar's test were employed for evaluation. Results show that Jackard-Needham and Dice metrics can display better performance for some descriptors. Yet, these performance differences were not found to be statistically significant.

arXiv:1512.02355v1 [cs.CV] 8 Dec 2015

Introduction: Binary image feature descriptors such as BRIEF [1], ORB [2] and BRISK [3] are becoming quite popular due to their easiness in computation, simple structure and high performance especially for applications aiming to run at decent video rates [4]. Many vision applications require finding image correspondences for computing a homography matrix which models the perspective transformation between two images and represents a linear relationship between image features.

The process of calculating a homography involves finding matching features detected across two images and then using an established method, *e.g.* RANSAC, to find the transformation supported by the majority of the feature correspondences. Previous work has shown that spatial distribution of these image features across the images play an important role in the accuracy of the calculated homography [5].

In practice, matching binary features is performed using two approaches. The first approach is known as 'brute force' matching, conventionally done using the Hamming distance given that it can be implemented efficiently using *XOR* instruction on bit sets, in which matching is performed by comparing every descriptor in the first image with the descriptors from the second image. The second approach, known as Fast Library for Approximate Nearest Neighbours (FLANN) [4], makes use of either a randomized kd-tree or a hierarchical clustering tree, depending on the feature dataset, to optimize performance for speed.

This paper investigates use of alternative approaches which are already used for comparing binary sequences in machine learning and pattern recognition [6] for the former matching method. A statistical evaluation will show the performance differences, in terms of the accuracy of the calculated homography matrix, when various distance metrics are employed.

Distance metrics: Considering two binary descriptors to be matched as two binary sequences, the following four dependent quantities are defined as:

f_{00} : num. of positions where both descriptors have 0s.
f_{01} : num. of positions where the first has 0 and the second has a 1.
f_{10} : num. of positions where the first has 1 and the second has a 0.
f_{11} : num. of positions where both descriptors have 1s. (1)

Using these quantities, the following distance metrics used for evaluation are defined: d_H for Hamming distance, d_J for Jaccard-Needham, d_C for correlation, d_D for Dice and d_Y for Yule distance metric. These are given by the literature [6, 7] as follows:

$$\begin{aligned} d_H &= \frac{f_{11}+f_{00}}{f_{00}+f_{01}+f_{10}+f_{11}} \\ d_J &= \frac{f_{10}+f_{01}}{f_{11}+f_{10}+f_{01}} \\ d_C &= \frac{1}{2} - \frac{f_{11}f_{00}-f_{10}f_{01}}{2\sigma} \\ d_D &= \frac{f_{10}+f_{01}}{2f_{11}+f_{10}+f_{01}} \\ d_Y &= \frac{f_{10}f_{01}}{f_{11}f_{00}+f_{10}f_{01}} \end{aligned} \quad (2)$$

where $\sigma = \sqrt{(f_{10}+f_{11})(f_{01}+f_{00})(f_{11}+f_{01})(f_{00}+f_{10})}$. These various metrics, widely used in data mining applications, are employed here in order to assess the accuracy of the homography matrix calculated using them.

Evaluation: The main question of this evaluation is first modelled in a *null hypothesis* framework (two-way ANOVA with image pairs and distance metrics as the two independent variables). Here, the null hypothesis states that there are no differences between the accuracies of the homography matrices calculated using the correspondences found when various metrics are employed, given the same set of keypoints. The alternative hypothesis states that there will be differences if different metrics are used for matching these binary descriptors. Further analysis is performed using McNemar's test to identify pairwise differences.

In order to quantify the accuracy of the calculated homography matrices, this evaluation employed the following approach depicted in Fig. 1: First image features are extracted using the detector part of ORB (5000 keypoints were extracted in order to provide good coverage [5]). These keypoints, with computed descriptors (*e.g.* BRIEF, BRISK), are matched across a pair of images. Using these matches, a homography matrix is calculated and applied to the first image in order to warp it onto the second. Here, an intermediate image is calculated ($d1$) which is then subtracted from the second image in order to remove the non-overlapping part ($d2$). Finally, the result image ($d3$) is obtained as the difference between warped version of the first image and $d2$. The evaluation criterion is chosen as the sum of non-zero pixels of $d3$. A larger sum indicates accuracy problems in the homography matrix resulting in alignment issues.

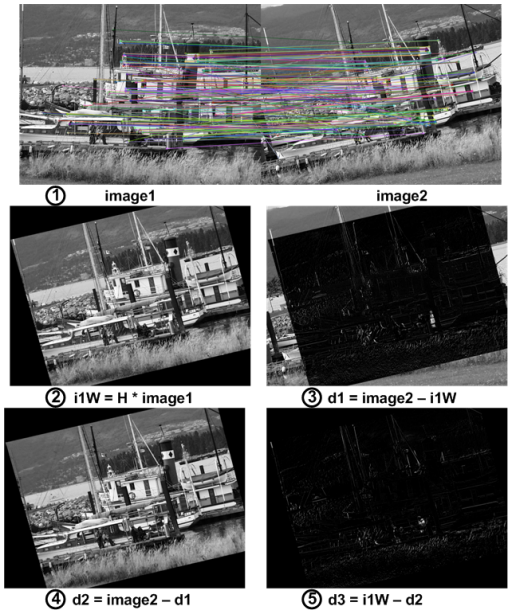

Fig. 1. Process used in the evaluation

It is important to reiterate that the evaluation here focuses only on the selection of the distance metric for finding a correspondence, not for evaluating different descriptors. However, descriptors were varied in type and size (*e.g.* BRIEF 32 and 64 bits) in order to see how this was affecting the results.

Using the sums of remaining pixels (samples for this evaluation) obtained from a dataset of 288 image pairs (all pairwise combinations) from a publicly available dataset[1], we applied a logarithmic transformation to the samples in order to reduce variance.

Results: Initial results from the samples (Fig. 2) show some differences across the various metrics and descriptor types. In order to see whether these differences are statistically significant or not, ANOVA was employed (Table 1). Analysis showed that different metrics have an effect on the results and the test confirms that there are significant differences ($F \gg Fcrit$, for all descriptor types). Furthermore, the *P-value* shows that the confidence on these results is very high.

Having found the differences, the next step is to identify the distance metric which is producing better results in terms of the accuracy of calculated homography matrix. McNemar's test [8] was employed for this purpose allowing pairwise comparisons between different metrics using z-scores. A value of 0 indicates no difference in performance, while 2.576

[1] `http://www.robots.ox.ac.uk/~vgg/research/affine/`

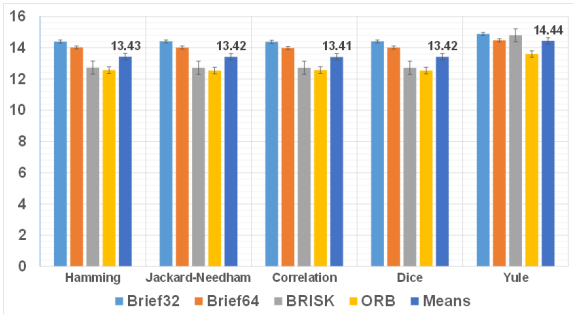

Fig. 2 *Effect of distance metric for different feature descriptors. Means bar indicate the mean performance of the metric in various descriptors.*

Table 1: Results for the ANOVA test

Source of Variation	*df*	*MS*	*F*	*P-value*	*F crit*
		BRIEF32			
Image pairs	287	26.26	88.83	0.00	1.16
Metrics	4	13.75	46.53	0.00	2.38
Error	1148	0.30			
		BRIEF64			
Image pairs	287	42.51	93.07	0.00	1.16
Metrics	4	12.52	27.41	0.00	2.38
Error	1148	0.46			
		BRISK			
Image pairs	287	134.53	63.60	0.00	1.16
Metrics	4	248.40	117.44	0.00	2.38
Error	1148	2.12			
		ORB			
Image pairs	287	154.47	128.49	0.00	1.16
Metrics	4	62.36	51.87	0.00	2.38
Error	1148	1.20			

indicates 99.5% confidence in one-tailed prediction meaning that one metric is statistically better than the other [8]. Looking at the results of Table 2, it can be seen that Yule displayed the worst performance of all the metrics for different descriptors. For BRISK and ORB descriptors, Jaccard-Needham has surpassed the Hamming distance; however, the difference was not statistically significant. Comparing this result with that of Fig. 2, on average, there are metrics that perform better than the Hamming distance, but this did not manifest itself significantly in a statistical evaluation.

Table 2: Results for the McNemar's test. Arrow-heads point to the metric resulting in a better accuracy.

Metric	Jaccard-Needham	Correlation	Dice	Yule
		BRIEF32		
Hamming	← 0.84	← 0.71	← 0.84	← 7.60
Jaccard-Needham		← 0.07	0.00	← 7.72
Correlation			↑ 0.07	← 7.72
Dice				← 7.72
		BRIEF64		
Hamming	← 1.30	← 0.45	← 1.30	← 4.30
Jaccard-Needham		0.00	0.00	← 4.42
Correlation			0.00	← 4.72
Dice				← 4.42
		BRISK		
Hamming	↑ 0.71	← 0.97	↑ 0.71	← 9.40
Jaccard-Needham		← 1.11	0.00	← 10.70
Correlation			↑ 1.11	← 10.11
Dice				← 10.70
		ORB		
Hamming	↑ 0.52	← 1.56	↑ 0.52	← 5.86
Jaccard-Needham		← 0.13	0.00	← 6.34
Correlation			↑ 0.13	← 5.98
Dice				← 6.34

Conclusion: This paper presented an evaluation in order to see whether distance metrics other than the conventional Hamming distance can be used for matching binary image feature descriptors. Results revealed that there are, indeed, metrics resulting in a more accurate homography matrix, producing less difference when the first image is warped onto the second. Findings also showed that these differences did not reflect well in a statistical evaluation, *i.e.* no significant differences were found. Considering the options provided by modern processor instruction sets for implementing the Hamming distance, it keeps its position to be the metric of choice for brute force matching of binary descriptors.

E. Bostanci (*Computer Engineering Department, Ankara University, Turkey*)

E-mail: ebostanci@ankara.edu.tr

References

1 Calonder, M., Lepetit, V., Ozuysal, M., Trzcinski, T., Strecha, C. and Fua, P.: 'BRIEF: Computing a Local Binary Descriptor Very Fast', *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2012, **34**, (7) pp. 1281–1298

2 Rublee, E. and Rabaud, V. and Konolige, K. and Bradski, G.: 'ORB: An efficient alternative to SIFT or SURF', *International Conference on Computer Vision*, 2011, pp. 2564-2571

3 Leutenegger, S., Chli, M. and Siegwart, R.Y.: 'BRISK: Binary Robust invariant scalable keypoints', *International Conference on Computer Vision*, 2011, pp. 2548-2555

4 Muja, M. and Lowe, D. G.: 'Fast Matching of Binary Features', *Ninth Conference on Computer and Robot Vision*, 2012, pp. 404–410

5 Bostanci, E., Kanwal, N. and Clark, A. F.: 'Spatial Statistics of Image Features for Performance Comparison', *IEEE Transactions on Image Processing*, 2014, **23**, (1), pp. 153–162

6 Choi, S. S., Cha, S. H. and Tappert, C.: 'A Survey of Binary Similarity and Distance Measures', *Journal on Systemics, Cybernetics and Informatics*, 2010, **8**, (81), pp. 43–48

7 Warrens, M. J.: 'Similarity Coefficients for Binary Data', *PhD Thesis*, Universiteit Leiden, 2008

8 Bostanci, B. and Bostanci, E.: 'An Evaluation of Classification Algorithms Using Mc Nemar's Test', *Proceedings of Seventh International Conference on Bio-Inspired Computing: Theories and Applications*, 2013, **201**, pp. 15–26